

(12) UK Patent Application (19) GB (11) 2 370 438 (13) A

(43) Date of A Publication 26.06.2002

(21) Application No 0031423.7

(22) Date of Filing 22.12.2000

(71) Applicant(s)

Hewlett Packard Company
(Incorporated in USA - Delaware)
3000 Hanover Street, Palo Alto, California 94304,
United States of America

(72) Inventor(s)

Stephen Philip Cheatle

(74) Agent and/or Address for Service

Hewlett-Packard Limited
Intellectual Property Section, Filton Road,
Stoke Gifford, BRISTOL, BS34 8QZ, United Kingdom

(51) INT CL⁷

H04N 1/387

(52) UK CL (Edition T)

H4F FESK FGXD F31K

(56) Documents Cited

EP 0824246 A2

WO 99/09887 A1

(58) Field of Search

UK CL (Edition S) H4F FESK FESX FFE FGXD

INT CL⁷ H04N 1/387 1/393 5/262

Online databases: WPI, EPODOC, JAPIO, INSPEC.

(54) Abstract Title

Automated image cropping using selected compositional rules.

(57) An image processing device (1, Figure 1) for cropping an electronic image (4) comprises image processing means (6), and an image memory (8). The image is processed to identify one or more features that are relevant to the composition of the electronic image, each feature having compositionally significant properties determined from several predetermined properties. For example, considering Figure 2A, the building feature (11) has the properties that it is (a) centrally positioned in the image and (b) includes a relatively high level of detail. Corresponding compositional rules can be formulated for each of the properties; the device operates to select appropriate compositional rules for the relevant image feature, and uses these to crop the image automatically around this feature. As shown in Figures 2B-2F, by analysing the image as grid regions (41), areas of the image having the similar colour/intensity are grown and merged, and the resulting blurred image is analysed, allowing cropping around areas of interest (Figure 2G).




Fig. 2F

BEST AVAILABLE COPY

1/10




Fig. 1



Fig. 2A

2/10




Fig. 2B




Fig. 2C




Fig. 2D

3/10

47 50




Fig. 2E




Fig. 2F

59




Fig. 2G

4/10

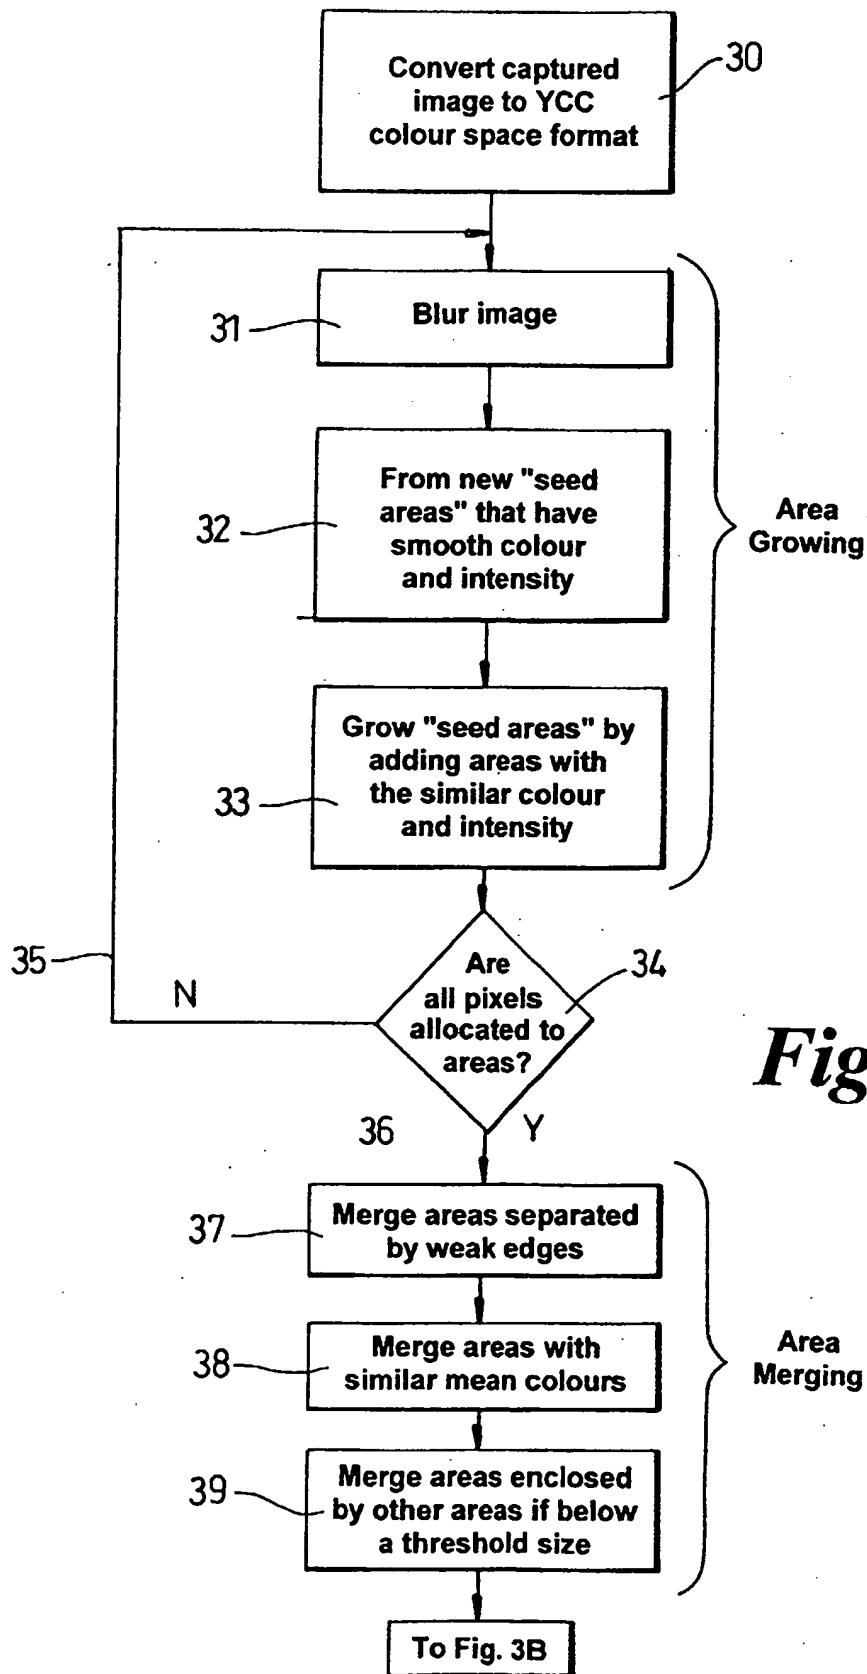


Fig. 3A

5/10

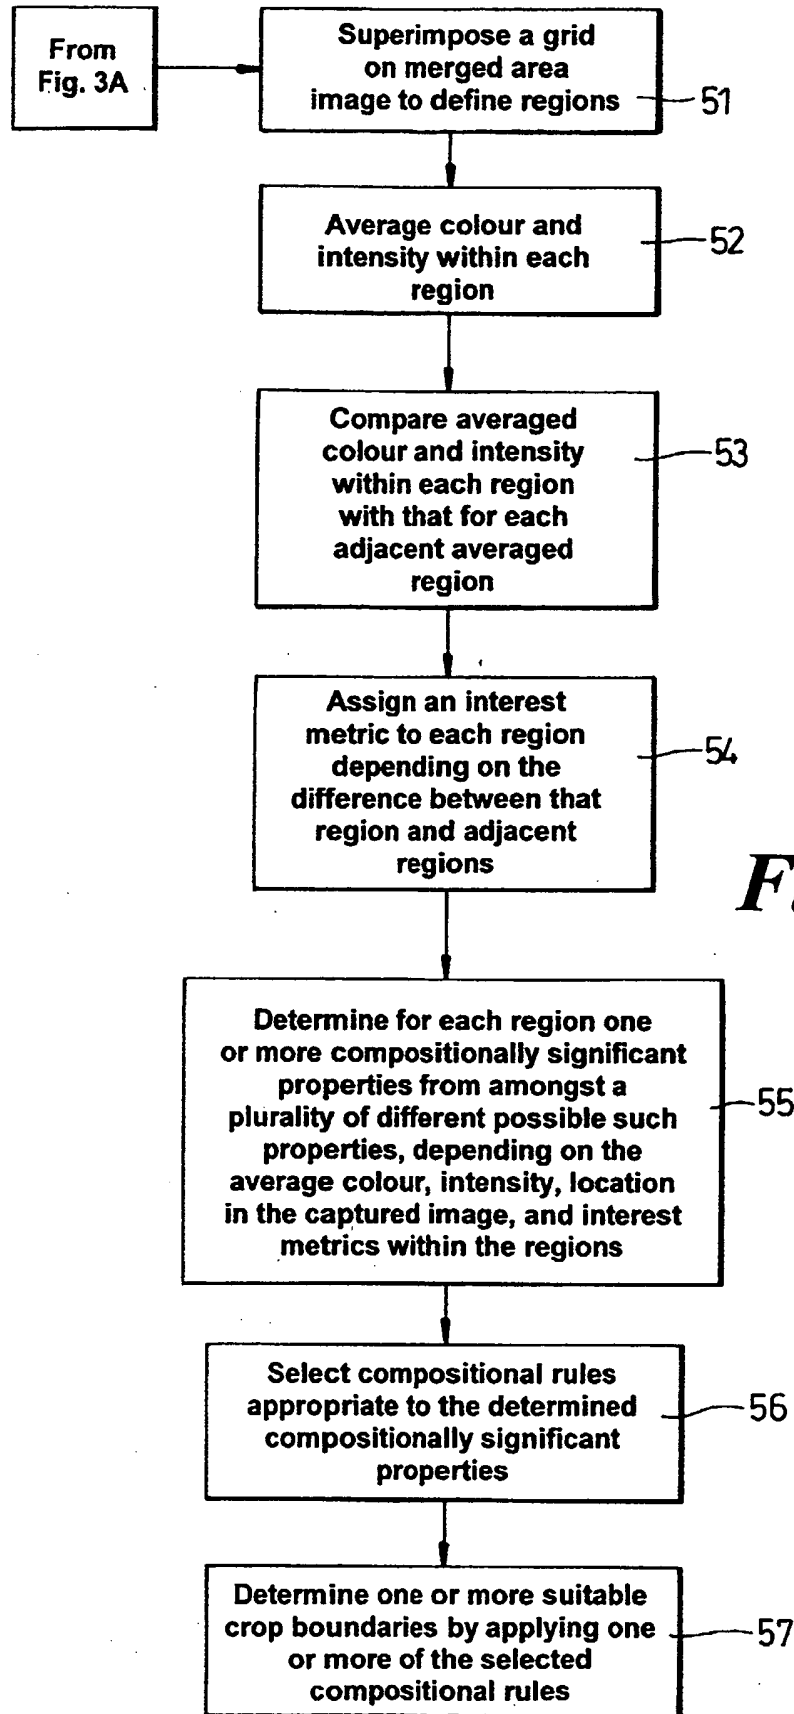


Fig. 3B

6/10

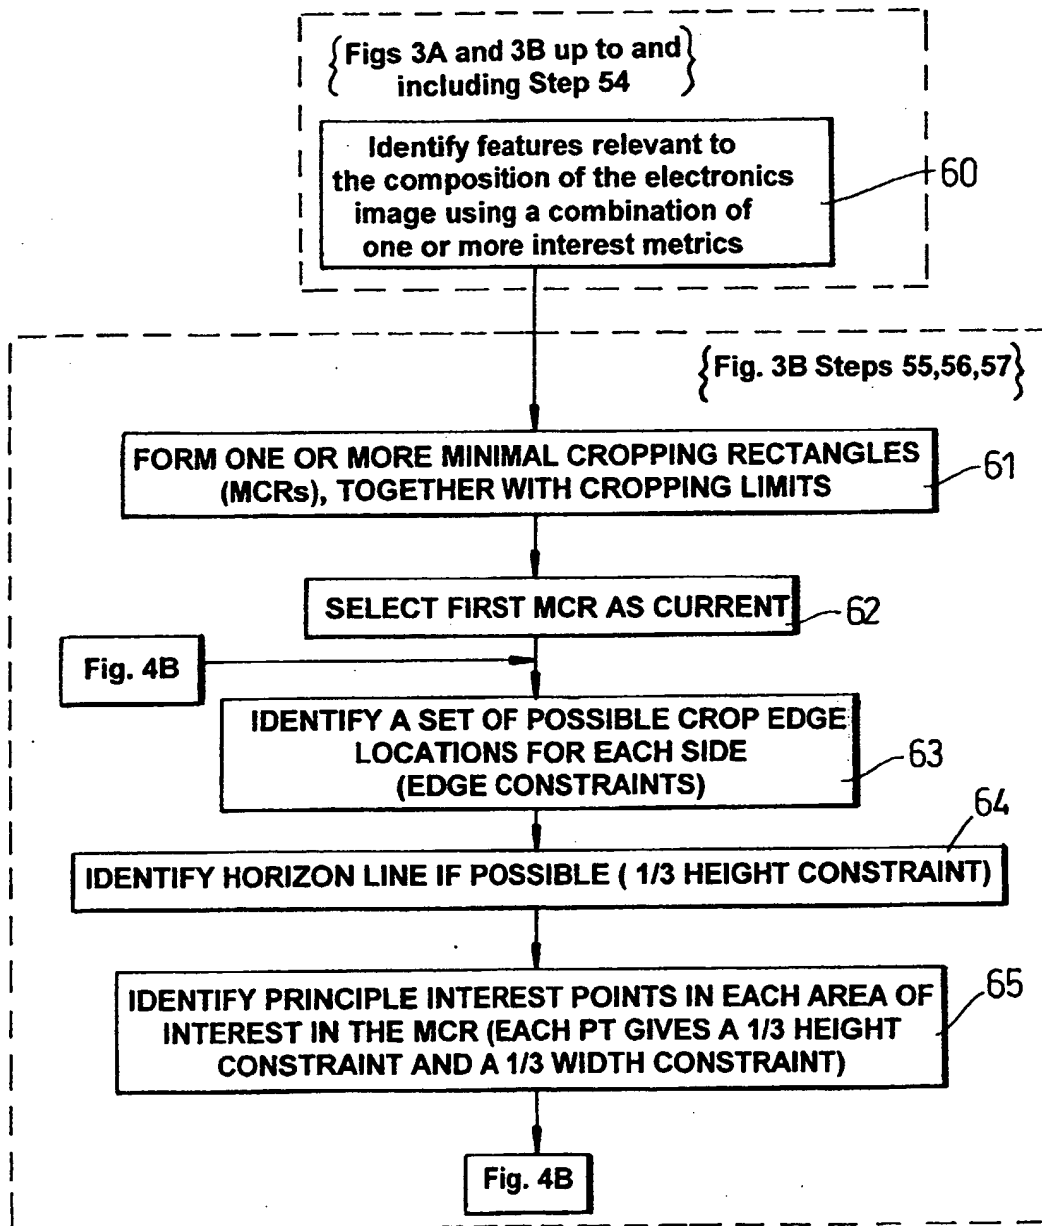
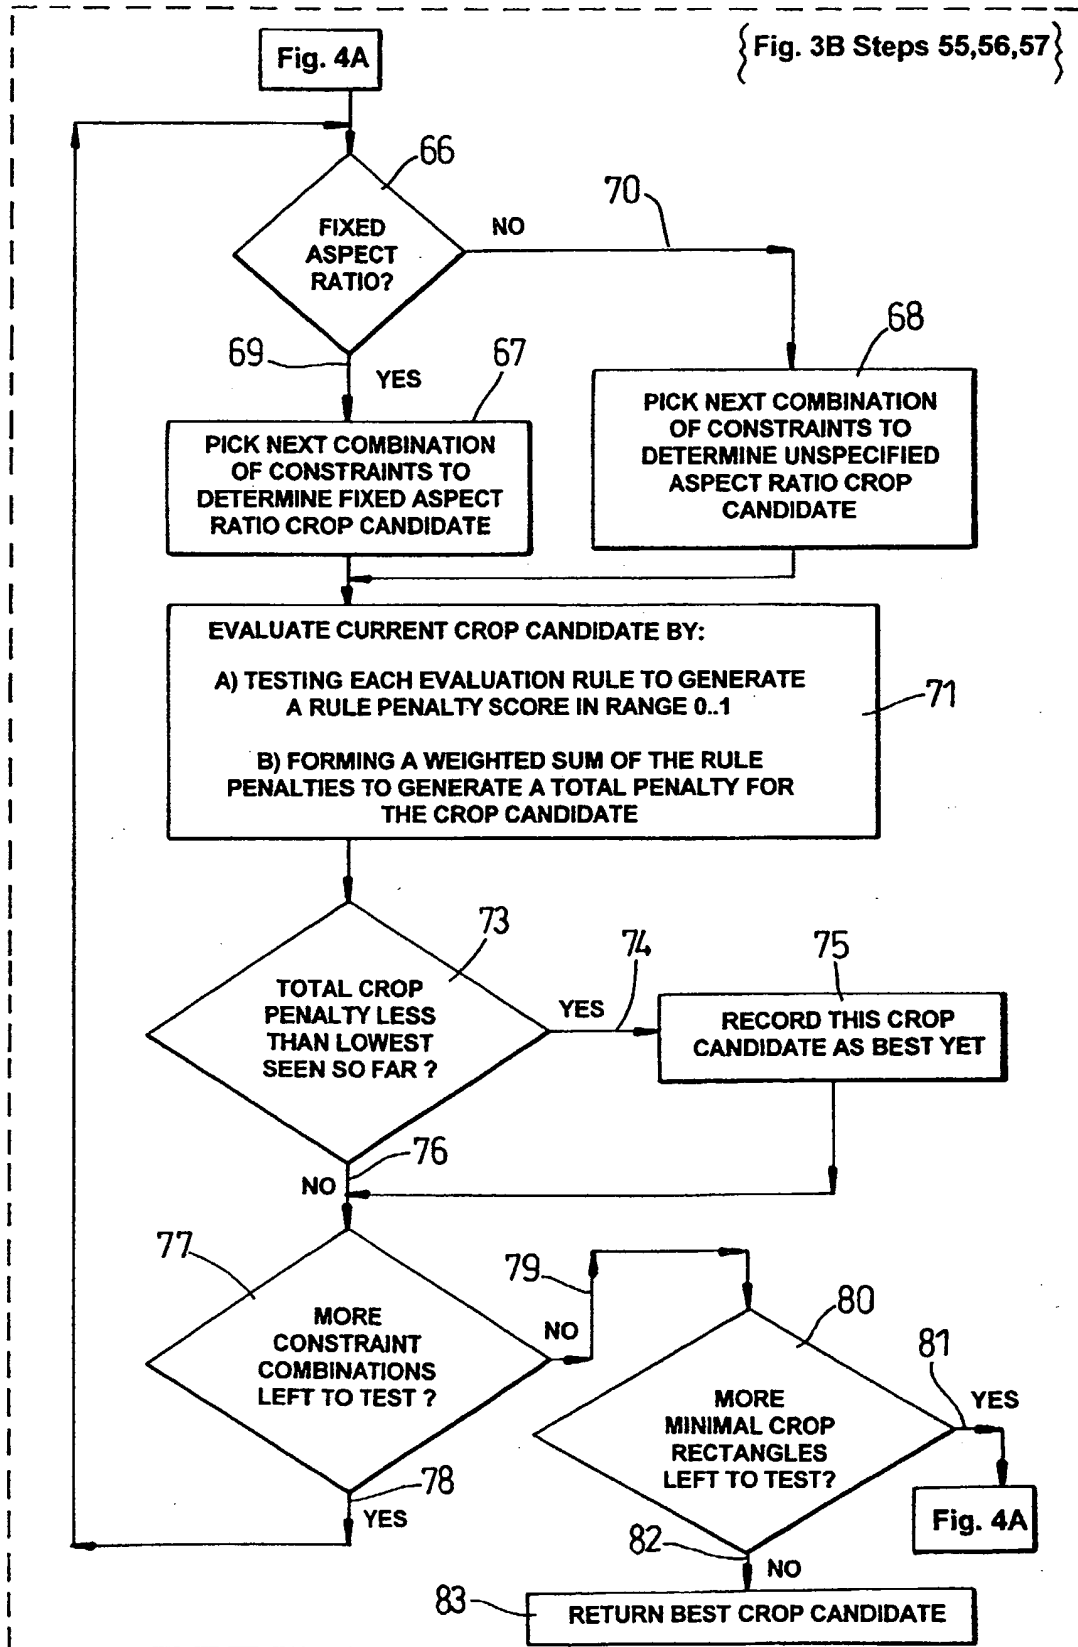


Fig. 4A

*Fig. 4B*

8/10




Fig. 5




Fig. 6

9/10

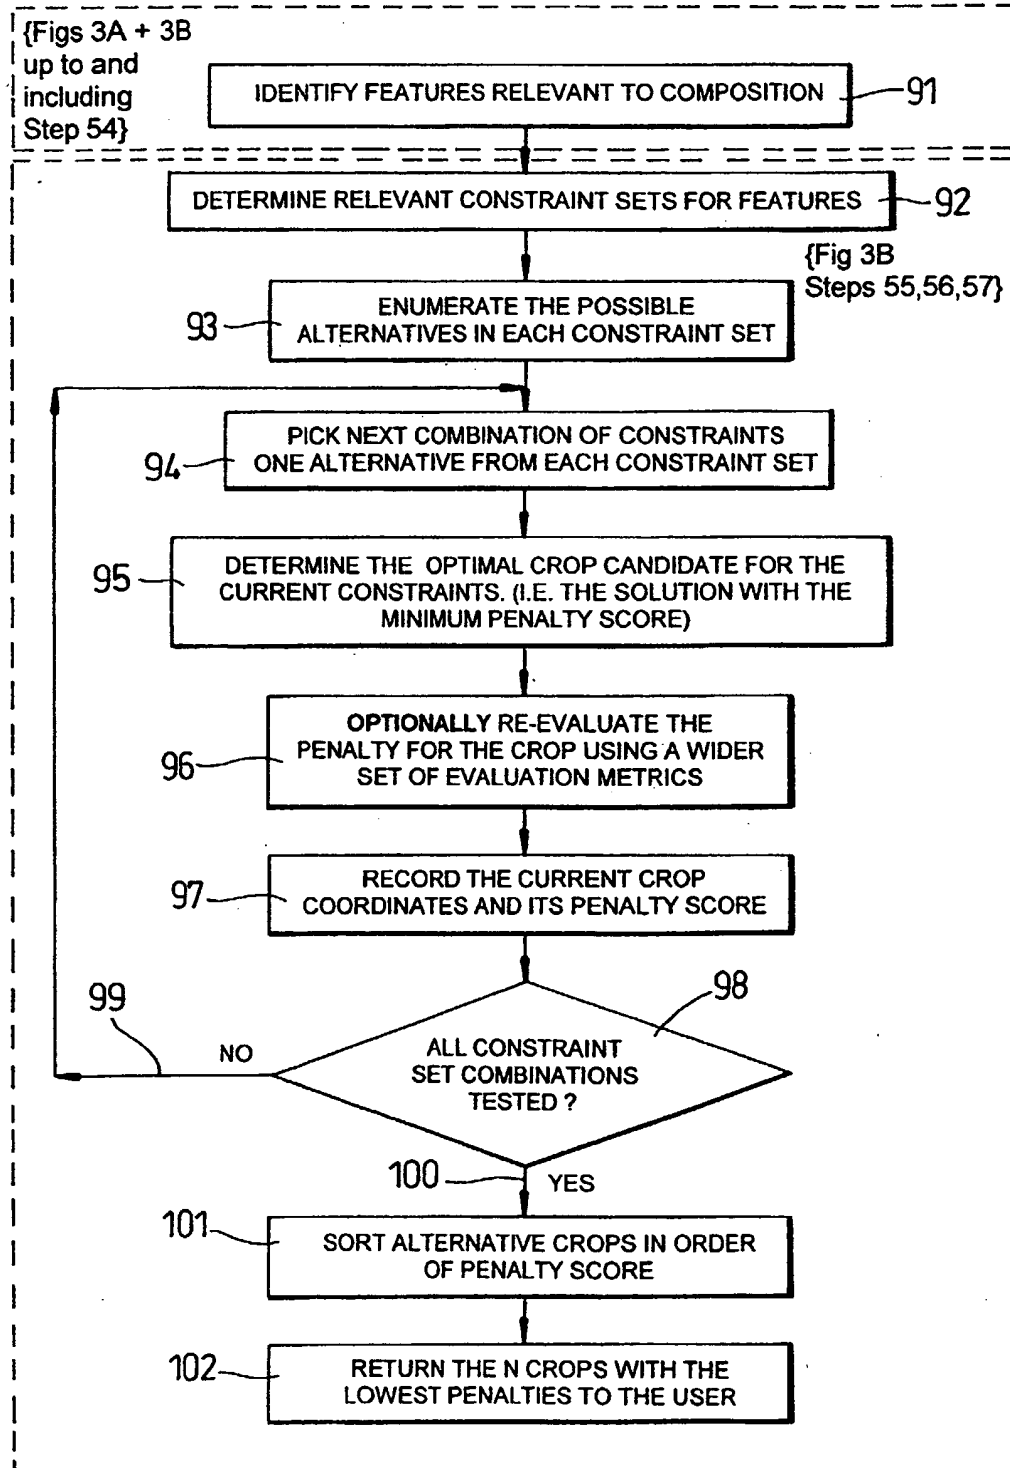


Fig. 7

10/10

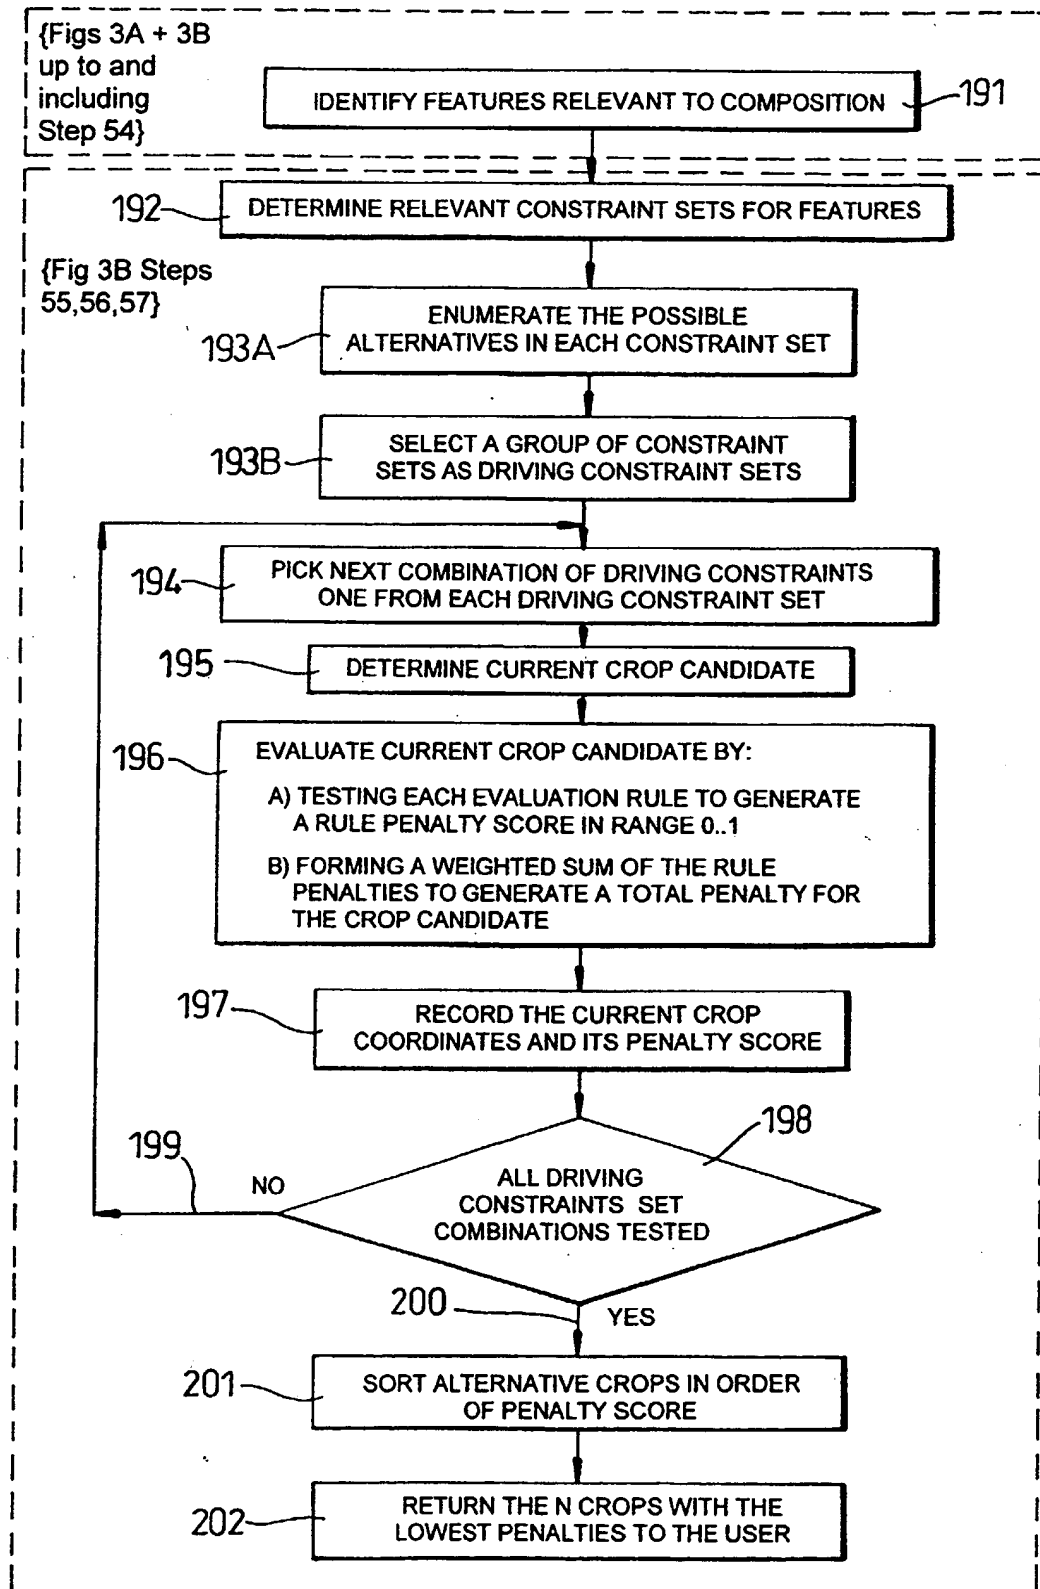


Fig. 8

Automated Cropping of Electronic Images

The present invention relates to the automated and semi-automated cropping of electronic images, and in particular to an apparatus and a method of using an electronic camera to capture and crop such electronic images.

Conventional photography requires the photographer to use a viewfinder both to aim the camera and to compose the picture. Composition at a particular location is done by changing the direction of the camera and altering the zoom control. Careful composition takes time and attention, as well as an understanding of various rules of good photographic composition. This is a skill that many find hard to learn. The effort required causes the photographer to be "out of the event" and in many cases this problem is enough to prevent potentially pleasing photographs from being taken. This is particularly the case when a photograph has to be taken quickly, for example when photographing action events, or children. Although, in principle, a photograph can be cropped after the event, this is time consuming and inconvenient, and may still require knowledge of the rules of good photographic composition, which the photographer may not possess.

It is an object of the present invention to provide a more convenient apparatus and method for capturing and cropping electronic images.

Accordingly, the invention provides an electronic image processing device for cropping an electronic image, comprising an image processing means, the image processing means including an electronic processor and firmware
5 and/or software for processing the electronic image, wherein the device is adapted to:

a) process the electronic image to identify one or more features relevant to the composition of the electronic
10 image, each such feature occupying a sub-region of the electronic image;

b) select one or more compositional rules from a set of predetermined compositional rules, based on the relevance
15 of the compositional rule(s) to one or more of the identified features; and

c) determine one or more suitable crop boundaries by applying one or more of the selected compositional rules.
20

Also according to the invention, there is provided a method of using an electronic image processing device for cropping an electronic image, the image processing device comprising an image processing means, the image processing
25 means including an electronic processor and firmware and/or software for processing the electronic image, wherein the method comprises the steps of using the image processing means to:

30 i) process the electronic image to identify one or more

features relevant to the composition of the electronic image each such feature occupying a sub-region of the electronic image;

- 5 ii) select one or more compositional rules from a set of predetermined compositional rules, based on the relevance of the compositional rule(s) to one or more of the identified features; and
- 10 iii) determine one or more suitable crop boundaries by applying one or more of the selected compositional rules.

In a preferred embodiment of the invention, each identified feature has one or more compositionally
15 significant properties from amongst a plurality of different predetermined compositional properties, and one or more compositional rules are selected from a set of predetermined compositional rules, based on the relevance of the compositional rule(s) to the compositionally
20 significant properties of one or more of the identified features.

The image may be cropped automatically by the image processing means according to a measure of the quality of
25 the potential crop boundaries. Alternatively, the user may manually select amongst the potential crop boundaries, so that the cropping of the image is performed semi-automatically. In this case, one or more of said suitable crop boundaries is presented to a user of the
30 device for manual selection by the user for cropping of

the electronic image.

Compositionally significant properties of a feature include things like the type of feature, e.g. a blank
5 area, an area of relatively high contrast or colour texture, recognisable objects such as a face, the sky or an edge.

The electronic imaging system may be part of an electronic
10 camera, or a document imaging system, or any other image capture system where the captured image may be cropped.

The electronic imaging system may be used with an electronic camera system for capturing an electronic image
15 of a scene. In particular, the electronic imaging system may be incorporated with the electronic camera.

Alternatively, the system may comprise a conventional electronic camera that outputs captured images to an image
20 processing device, for example a personal computer, that includes the image processing means.

The camera will, in general, comprise a detector array for capturing the electronic image and an optical imaging
25 system arranged to image the scene onto the detector array. The camera may be a hand-held still electronic camera and/or a video electronic camera.

The compositional rule may comprise edge placement
30 criteria, for example having a dark border to an edge of

the cropped image, or placing an edge one-third or two-thirds the way from an area of interest within the crop boundary.

- 5 It may be that just one compositional rule is used multiple times on a single image, once for each identified feature.

Once candidate crop boundaries have been determined, these
10 may be presented to a user of the system. The user may then choose a cropping candidate, whereupon the image processing device may be arranged to crop the electronic image according to the user's choice.

- 15 The identifiable features should be those features that are relevant to the placement of cropping boundaries in the processed image.

The camera may include means by which a user of the camera
20 may tag one or more features relevant to the composition of the electronic image, said tagged feature(s) then being associated with a compositional rule that includes said tagged feature(s) in the cropped image. Such tagging may be done by the user indicating by means of suitable
25 controls a feature or an area of the captured image as displayed to the user, for example on a LCD display built into the camera.

- 30 It may be possible, however, for features in the image to be tagged automatically. For example, a person in view of

the camera may wear some sort of identifier tag which can be recognised automatically by the image processing means within the camera. The tag may be an optically distinguishable badge pattern understood by the image processing software. A person can then be automatically identified.

Optionally therefore, the identifiable features may include a predetermined feature, for example a tag that a person may wear. Such a tag may have an identifiable pattern to the image processing system recognises. A compositional rule will then be associated with such an identified tag so that the identified feature is included in the cropped image. So that the tag does not distract a user viewing the cropped image, it is preferable if it is implemented as an infra-red only tag (for example, as an infra-red transmitter).

One way of manually tagging features is to use the camera viewfinder as a pointing device (as opposed to its conventional use as both a pointing and composition device). In this use the prime area of interest is deliberately positioned in the approximate centre of the frame. When the image is auto-cropped according to this invention, the region at the centre of the image is deemed to be essential to the cropped image and is thus prevented from being cropped out.

In another embodiment of the invention, prior to identifying in the captured image any features, the

electronic camera may be used to capture an image of an object with an appearance corresponding with a feature to be included in the cropped image. This may relate to the appearance of the object (or the person - for example, a face) itself, but in a different use model may also relate to appendages worn for the purpose of identification. For example, if a person is wearing a jacket of a particular shade of blue, then the electronic camera may be pointed at the jacket in order to "initialise" the image processing means to recognise that shade of blue as being associated with a tagged feature when that shade of blue is captured in an electronic image. This may be assigned a high interest metric and/or may be associated with a particular compositional rule. When an image is captured, the image processing means may then be used to identify in the captured image at least one tagged feature. Then, the compositional rules may be used to crop the captured image so that the tagged feature is included in the cropped image.

20

In one embodiment of the invention, step iii) comprises the steps of:

iv) generating a plurality of alternative candidate crop boundaries;

25

v) generating a measure of composition quality for each alternative candidate crop boundary by using a set of picture composition metrics; and

30

vi) selecting as output a reduced number of crop candidates having a relatively high measure of composition quality, for example, just one crop candidate.

5 In an alternative embodiment of the invention, step (iii) comprises the steps of:

vii) implementing the compositional rules as constraints that may be satisfied to a greater or lesser degree, each
10 constraint having an associated cost function that increases the less well the constraint is satisfied.

viii) defining an overall cost metric as a function of crop co-ordinates in the image, by forming a combination
15 of the separate cost functions associated with each individual constraint;

ix) applying an optimisation method to find one or more best crop locations by finding minima in the overall cost
20 metric; and

x) selecting as output a reduced number of crop candidates for said best crop locations, for example, just one crop candidate.

25

There will be at least one crop edge. For example a circle or an ellipse has just one crop edge. Often there will be more than one crop edge, for example a square or a rectangle has four edges, between one and four of which
30 will result from cropping of the original electronic

image.

When the crop boundary of the cropped image has a plurality of crop edges, the steps of the method may be performed separately for each crop edge to generate the crop boundary. This help to reduce the calculation needed to select the crop boundaries, as each edge may be assessed independently from the other edges.

One way in which the invention may be implemented is if step (i) described above includes the step of identifying features which constitute non-overlapping segmented regions of interest. Then step (iii) may be achieved by first selecting alternative divisions of said identified regions of interest into those which will be included by the crop boundary and those which will be excluded by the crop boundary. Each such alternative division of regions of interest is then used to determine an inner crop boundary limit and an outer crop boundary limit such that the inner crop boundary limit is the smallest boundary which circumscribes the regions of interest to be included and the outer crop boundary limit is the largest boundary which excludes the regions of interest to be excluded. The one or more suitable crop boundaries can then be determined between the limits of the inner and outer crop boundary limits for each of said alternative divisions of regions of interest.

When the crop boundary has a plurality of edges and one or more suitable crop boundaries are determined by evaluating

an edge quality metric for each possible edge placement between the inner and outer crop boundary limits, the selection of best crop boundary may at least in part be dependent on the separate edge quality metrics of each of the boundary's edges.

This can also help to reduce the computational burden, as it is then not necessary to consider alternative crops where the outer crop boundary limit does not completely enclose the inner crop boundary limit, or where the shape of the area between the outer and inner crop boundary limits does not permit the placement of the desired crop boundary shape.

It may be that in step (i) blank or uninteresting areas are detected as some of the features relevant to the composition and in steps (ii) and (iii) one or more of the compositional rules cause the image to be cropped according to the arrangement of said blank or uninteresting areas relative to other identified features relevant to the composition of the electronic image.

In photography, such blank areas are usually associated with a plain background, or sky. The image processing means can use one or more compositional rules associated with such blank areas. For example, a blank area extending across an upper horizontal band of the scene may be associated with plain sky, and so an appropriate compositional rule might be to minimise the amount of this particular area, and to orient horizontally the boundary

between this area and lower areas of interest. Therefore,
in general, the method may comprise the step of using the
compositional rules to crop the captured image according
to the arrangement of the blank areas relative to the
5 other features which are not blank.

If the image is a colour image, it may be that at least
one feature is identified by segmenting the image into
areas of similar colour and/or texture.

10

One way of performing the method according to the
invention is if features relevant to the composition of
the electronic image are identified in step (i) by a
process of: segmenting the image into regions; denoting
15 certain regions as more salient than others; grouping
salient regions into larger regions separated by
relatively non-salient regions; and identifying said
groups of salient regions as features with a region of
interest property.

20

Each such area of similar colour and/or texture may then
constitute one of the identified features.

The image may then be segmented into regions by segmenting
25 the image into homogenous sub-areas, the measure of
homogeneity being based on some combination of colour,
intensity and texture.

One way in which the image may be segmented is by first
30 dividing the captured image into a plurality of regions,

averaging the colour and/or the texture within each region. Adjacent similar regions may then be grouped to delineate at least one type of area of interest having a low interest metric. Additionally or alternatively,
5 adjacent dissimilar regions may then be grouped to delineate at least one type of area of interest having a high interest metric.

The term "interest metric" is used herein to define
10 weighting or importance attached to a particular area of interest. The interest metric for an identified feature may be used to associate with this feature at least one particular photographic compositional rule. For example a feature with a high interest metric may be associated with
15 a rule that places such an area either at the centre of a cropped image, or slightly to one side of the centre of an image, for example according to the well-known photographic compositional rule called "the rule of thirds".

20

The system, may also be capable of face detection, or person detection using one or more of: clothing recognition, hair colour and approximate style recognition, or use of body shape models.

25

Certain regions may be denoted as more salient than others by allocating high salience to regions on the basis of some combination of:

30 - relative unusualness of the colour, intensity or texture

of the region to other adjacent regions; and/or

- relative unusualness of the colour, intensity or texture
of the region relative to a substantial portion of the
5 image.

A decision may then be made by software and/or firmware as
to how likely the region is to be a person's face, head or
whole body, or how likely the region is to be a known
10 person or how central the region is in the image.

The system may also allow for interactive input from the
user to indicate the prime region of interest.

15 The compositional rule for features identified at the
centre of the image may be to include such features in the
cropped image.

The invention will now be described by way of example
20 only, with reference to the accompanying drawings, in
which:

Figure 1 is a block schematic view of an electronic camera
system incorporating an electronic image processing device
25 according to the invention, having a detector array and an
image processor for automatically cropping electronic
images captured by the detector array;

Figures 2A to 2G are illustrations showing how a poorly
30 composed photograph can be automatically cropped;

Figures 3A and 3B are flow charts illustrating an embodiment of a method according to the invention for cropping a captured electronic image;

5

Figures 4A and 4B are flow charts showing in detail one way of implementing the method according to the invention as shown in Figures 3A and 3B;

10 Figure 5 is a plot of crop penalty metric against crop distance for alternative positions of a main region of interest;

Figure 6 shows schematically minimum and maximum cropping rectangles for two identified features;

15

Figure 7 is a flow chart showing in detail one way of implementing the method according to the invention as shown in Figures 3A and 3B, using a "constraint" based approach; and

20

Figure 8 is a flow chart showing in detail one way of implementing the method according to the invention as shown in Figures 3A and 3B, using a "generate and test" based approach.

25

Figure 1 shows schematically an electronic camera 1 for capturing an electronic image of a scene 2. The camera has an optical imaging system 3 arranged to image the scene 2 onto a two-dimensional detector array 4 within the camera.

30

The detector array is connected 5 to a microprocessor 6 for processing of the images captured by the detector array. The microprocessor 6 has internal software and/or firmware for processing captured electronic images.

5

The microprocessor 6 is itself connected to an image capture button 7 by which a user of the camera 1 may initiate capture of the image, and also to a memory 8. Not shown are various items normally associated with a
10 conventional electronic camera, namely a battery power supply, viewfinder or liquid crystal viewfinder display, focus and light level detection optics and electronics, and exposure and auto-focus control mechanisms.

15 Figure 2A shows an electronic image 10 captured with the electronic camera 1. The image 10 has as its central feature a farmhouse 11. Other features include: a dark tree 12, to the left of the farmhouse; a horizon line 13 extending across a right-hand portion of the image 10; and
20 above the horizon line 13, a generally cloudy sky 14. The foreground consists of another feature, namely a nearly featureless field 15. There are various distracting details 16 in a left-hand portion of the image 10 that form one or more other features.

25

All of these features occupy various sub-regions or areas of the captured image 10, and all are relevant to the composition of the electronic image, having various compositionally significant properties. For example,
30 compositional properties of the farmhouse 11 include its

central location in the captured image, and the relatively high level of detail compared with other features, making it a main area of interest. The sky 14 and field 15 share the compositional property that these extend to borders of the captured image, and are relatively featureless. The horizon line 13 has the compositional property that it is non-central, horizontal and well-defined in terms of contrast. The details 16 have the compositional property that they are not central to the captured image 10, and that they are a competing area of interest to the main area of interest, i.e. the farmhouse 11.

Compositional rules can be formed for each of these compositional properties.

15

According to conventional rules of photographic composition, the image 10 is poorly composed. In particular, the main area of interest to the eye is the farmhouse 11, which is dwarfed by the expanse of sky 14 and field 15. The extended horizon line 13 and details 16 detract from the main subject of the farmhouse 11.

Figures 2B-2G show one way in which the poorly composed image 10 can be automatically cropped to leave a much more satisfactorily composed image 59, as shown in Figure 2G. The first step is to process the image 10 in such a way that areas with similar colour, intensity and texture are merged into areas having a uniform colour and intensity, as shown in the image 20 in Figure 2B.

30

An example of a method for producing the image 20 with merged areas is shown in Figure 3A. First, the captured image 10 is converted 30 into an image having a compressed colour variation while still retaining intensity variations, such as an image converted to the YCC colour space format. Areas within the YCC format image having similar colour and intensity are then generated and grown, by first blurring the image 31 then identifying in the blurred image "seed areas" that have smooth or uniform colour and intensity 32. Seed areas are then grown by adding areas to the boundary of seed areas that have sufficiently similar colour and intensity 33. A test 34 is then performed as to whether or not all pixels in the YCC format image have been allocated to grown seed areas. If not 35, then the original YCC format images are blurred to a greater extent 31, following which seed areas are again identified 32 and grown 33 as described above. At some point, the blurring 31 of the image will be sufficient such that all pixels are allocated to grown seed areas 36.

20

The processed image is then further simplified by merging adjacent areas separated by weak edges 37. Weak edges are those separating areas having relatively low colour or intensity differences. Next, areas with similar mean colours are merged 38. Finally, small areas enclosed by other larger areas are merged with the larger areas if the small areas fall below a threshold size 39. When this process is applied to the captured image 10, the result is a much simplified image 20 having a relatively small number of merged areas five of which 21-25 are indicated

30

in Figure 2B. Each of these areas 21-25 has a uniform colour and intensity.

5 The rest of the process for automatically cropping the image is shown in Figures 2C-2G and illustrated in the flow-chart of Figure 3B. First, an analysis grid 40 is superimposed 51 on the processed image 20 having the merged areas 21-25. The grid 40 defines a number of square regions 41 and in each of these the colour and intensity
10 is averaged 52. The average colour and intensity within each region is compared with that for each adjacent averaged region 53, and the differences are used to assign an "interest metric" to each region 41 depending on the difference between that region and adjacent regions 54.
15 The various interest metrics are shown in Figure 2D as squares 45 of the same size as the regions 41 and having a brightness proportional to the difference in averaged colour and intensity with adjacent regions 41.

20 It should be noted that the description up to this point is just one example of how the initial processing of the image may be performed to identify different features relevant to the composition of the electronic image.

25 Each merged area 21-25 can then be classified 55 according to various criteria, for example its average colour and intensity, location in the captured image 10 and the various interest metrics 45 within each merged area 21-25. From this, it is possible to determine for each region one
30 or more compositionally significant properties from

amongst a plurality of different possible such properties. For example, the merged area 21 for the field 15 in front of the farmhouse 11 has very low interest metrics 45, as does the merged area 22 for the sky 14 above the farmhouse 11. The farmhouse 11 results in an area with high interest metrics 45, and the areas of distracting detail 16 to the left of the farmhouse are assigned moderate interest metrics.

10 Areas with low interest metrics 45 extending fully across the top and bottom of the captured image 10 can be associated with a compositional rule which allows these areas 21, 22 to be substantially cropped. Merged areas 23 which result in a local maximum in the interest metrics 15 separated from the highest interest metrics 45 can be identified as competing areas of interest 16 that may detract from the main area of interest 24. Again, compositional rules can be associated with these merged competing areas of interest 23 that allow these to be 20 cropped.

A simple compositional rule, such as the "rule of thirds" can be associated with the main area of interest, here the farmhouse 11. A dark merged area 25 from the tree 12 to 25 the left of the farmhouse 11 can be associated with the compositional rule that it is desirable to use dark areas away from a main area of interest as a natural boarder to the cropped image. Therefore, various compositional rules appropriate to each type of merged area 21-25 can be 30 selected 56. This is shown in Figure 2E, where a three-by-

three rectangular grid 50 is superimposed over the grid of areas of interest 45 in such a way that the left-hand edge 47 coincides with the location of the dark merged area 25 associated with the dark tree 12, and with a one-third line 42 running vertically through the highest interest metrics 45. The rectangular three-by-three grid 50 is then positioned vertically so that a highest interest metric 45 is centred vertically in the grid 50. Therefore, by applying the selected compositional rules, a suitable crop boundary can be determined 57. The system may then present this crop boundary to a user of the system as a suggested crop boundary.

Figure 2F shows a captured image 10 with the rectangular three-by-three grid 50 positioned as determined according to the procedure described above. The captured image 10 can then be cropped along the outside boundaries of the rectangular grid 50 to result in an image 59 cropped according to the selected compositional rules. It will be apparent that the cropped image 59 has a better composition than the original captured image 10.

Preferably, the processing of the captured image 10 is performed by the microprocessor 6 in the camera 1, with the memory 8 serving to hold data generated by the automatic cropping process, and the resultant cropped images 59. It would, however, be possible for image processing to be done external to a camera body, in which case the electronic camera and external processing form the electronic camera system of the invention.

The invention is particularly useful in the case of electronic cameras having a detector array 4 with a relatively large number of detector elements. For example,
5 a detector array having two million or more elements can be used with an optical system 3 having a wide angle field of view. The user then need only point the electronic camera 1 generally in the direction of a scene 2 he wishes to photograph. Automatic cropping can then be used as
10 described above to crop unwanted areas of the captured image. This relieves a considerable burden from the photographer, as he no longer needs to worry unduly about details of photographic composition. Electronic photographs can then be taken rapidly, which increases the
15 chances that the photographer will capture the desired moment.

One way of manually tagging features is to use the camera viewfinder as a pointing device (as opposed to its
20 conventional use as both a pointing and composition device). In this use the prime area of interest is deliberately positioned in the approximate centre of the frame. When the image is auto-cropped according to this invention, the region at the centre of the image is deemed
25 to be essential to the cropped image and is thus prevented from being cropped out.

In another variant of the invention, the camera may be initialised to identify certain colours or texture as
30 having a high interest. At least two use models are

possible here. One simply involves the identification of features of "natural" importance or inherent interest: faces, the overall shape of a person or object, and other expected compositional elements. Another is to provide
5 additional elements or appendages for the specific purpose of "tagging" to force inclusion of an object in the cropped image. The practical effect may be similar in either case. For example, if a person is wearing a blue rain jacket, then the camera may be pointed close up at
10 the blue rain jacket and then capture an image of the jacket. The camera can then be programmed to process a captured image to assign that particular colour a high interest metric. If a wide angle picture is then taken of a scene in which the blue jacket appears, then this area
15 can be assigned the highest interest metric so that the captured image is automatically cropped in such a way that the blue jacket is retained in the image. This is particularly useful when images are captured of a crowd of people, one of which the photographer would like to make
20 the main subject of the photograph.

Figures 4A, 4B and 5 show in detail possible steps for implementing the steps 55, 56 and 57 briefly outlined above in Figures 3A and 3B.

25

Once features relevant to the composition of the electronic image have been identified 60, as the steps up to and including step 54, a minimal cropping rectangle, (MCR), and associated cropping limits can be defined 61 as
30 crude limits on the minimum and maximum areas of a cropped

image. If the interest map for an image generated in 60 shows N distinct areas of interest, (for example, areas of interest separated by some area of non-interest determined by some adaptively set threshold), possible minimal
5 cropping rectangles can be generated which contain alternative combinations of between 1 and N areas of interest where the MCR contains the selected combination of areas of interest and excludes the other areas. Note that not all combinations may be possible, as these may
10 not be contained within a single rectangle that excludes the non-selected areas. The cropping limit for a given MCR is the maximum rectangle, which contains the MCR but excludes the non-selected areas of interest. This is shown in Figure 6, which shows for a captured image 90 two
15 different minimum cropping rectangles 84,85, that will be explained in more detail below.

Each MCR 84 or 85 and its associated cropping limit is processed in turn. The first step is to select a first MCR
20 84 or 85 as a current potential cropping candidate 62. Steps 63 to 70 essentially set up a set of preferred constraints on possible crop locations and then successively generate candidate crops that meet some combination of these constraints. In practice there will
25 typically be too many constraints to satisfy all constraints simultaneously. Step 63 in particular, in which a set of possible crop edge locations is identified for each location as edge constraints, deliberately suggests multiple possible locations for each edge.
30 However, only one possibility can be used in any one crop.

Steps 64 and 65 are given as examples of the implementation of the compositional "rule of thirds" as applied to the identified features of horizon line 64 and
5 a main area of interest 65.

The idea behind Figures 4A and 4B is that all combinations of position are generated and then evaluated in the later steps. This is a "generate and test" approach to
10 determining one or more suitable crop boundaries, using a "constraint set" determined by minimum and maximum cropping rectangles.

Note that step 66, where a question is asked as to whether
15 or not the cropped image is to have a fixed aspect ratio, ensures that it is possible that a particular aspect ratio can be specified and enforced. In practice this means that when an aspect ratio is specified 69, a smaller number of other constraints will in general be required 67 to
20 completely specify a crop candidate, than in the alternative case 68 when no aspect ratio is required 70. In the case of no explicit aspect ratio requirement, it is likely that an evaluation rule in step 10 will penalize ridiculously thin aspect ratios.

25

Once a crop candidate has been identified, it is then evaluated 71 by one or more rules. Each rule is implemented as a heuristically evaluated measure on the image. For example, a metric 72 which measures how close a
30 point of interest is from a one-third line is shown in

Figure 5. The fractional position of the point in the candidate crop is measured in both the horizontal and vertical directions. The penalty for each direction is determined from the heuristically determined graph shown in Figure 5. The two measures $penalty_{vert}$ and $penalty_{horiz}$ are combined by the rule:

```
penalty = max(  $penalty_{vert}$  ,  $penalty_{horiz}$  )  
if max(  $penalty_{vert}$  ,  $penalty_{horiz}$  ) > 0.75
```

10

and

```
penalty = mean(  $penalty_{vert}$  ,  $penalty_{horiz}$  )  
if max(  $penalty_{vert}$  ,  $penalty_{horiz}$  ) ≤ 0.75
```

15

Similar heuristic measures are used for other compositional rules such as eliminating distractions close to the edge of the frame, minimum edge quality, a preference for dark or low activity boundaries, and so on.

20

The combination of different rule penalties by a weighted sum allows some rules to be considered as more important than others. Again, the weightings are determined heuristically.

25

There are many possible extensions of the basic scheme. For example, it would be possible for the rule combination weightings to be dynamically adjusted according to the overall type of image. For example, a MCR 84 or 85 with a single area of interest containing a single face looking

30

straight towards the camera may reduce the weighting for the rule of thirds, allowing a more centrally placed portrait to be preferred.

- 5 Another possibility is for an additional penalty factor to be generated from step 61 where some MCRs are intrinsically preferred (i.e. given a low penalty) compared to others.
- 10 In Figure 4B, the penalty is evaluated as follows. First, a test 73 is performed as to whether or not the total crop penalty is less than a lowest previous total crop penalty. If so 74, then the current crop candidate is recorded 75 as the best crop candidate so far. If not 76, then a test
15 77 is performed as to whether or not there are more constraint combinations left to test. If so 78, then the flowchart loop back to step 66.
- If not 79, the flow chart next tests 80 if there are other
20 minimal cropping rectangles 84 or 85 left to test. If so 81, then the flow chart loops back to step 63. If not 82, the flow chart shows that the best crop candidate is returned 83 as an output from the process.
- 25 Figures 7 and 8 show in more general terms the concepts behind respectively the "constraint-based" and "generate and test" approaches described above. In both cases, the starting point is the identification of features relevant to the composition of the image 91,191, for example as set
30 out in the steps up to an including step 54 in Figures 3A

and 3B.

The next step is to determine 92 the relevant "constraint set" for the identified features. The concept of a
5 "constraint set" is a set of alternatives, only one of which should be considered at a time. Each alternative consists of one or more fully specified constraints, for example features with some required value, which are then enumerated 93.

10

A simple example of a constraint set is "aspect ratio". There are two alternatives, "portrait" and "landscape". Alternative 1 (portrait) might be defined by the constraint:

15

$$(Right - Left) / (Bottom - Top) = 0.75$$

Alternative 2 (landscape) might be defined by the constraint:

20

$$(Right - Left) / (Bottom - Top) = 1.33$$

With reference to Figure 6, a more complex constraint set could define cropping limit alternatives for different
25 groupings of areas of interest 84,85 within the maximum boundary 90 of the original captured image. The number of alternatives in this set is determined by the analysis of the areas of interest in the image. Suppose that two areas of interest 84,85 have been determined with pixel
30 co-ordinates: (30,180) to (200, 300) and (350,50) to

(600,240) as shown below. In this example, the whole image has a top left co-ordinate (0,0) and bottom right co-ordinate (640,480).

- 5 The crop limits constraint set would consist of three alternatives:

Alternative 1 (left-hand area of interest 84 only) is defined by the constraints:

10

Left > 0 *Left* < 30
Top > 0 *Top* < 180
Right > 200 *Right* < 350
Bottom > 300 *Bottom* < 480

15

Alternative 2 (right-hand area of interest 85 only) is defined by the constraints:

20 *Left* > 200 *Left* < 350
Top > 0 *Top* < 50
Right > 600 *Right* < 640
Bottom > 240 *Bottom* < 480

- 25 Alternative 3 (include both areas of interest 86) is defined by the constraints:

30 *Left* > 0 *Left* < 30
Top > 0 *Top* < 50
Right > 600 *Right* < 640
Bottom > 300 *Bottom* < 480

The constraint set concept can be used to represent many mutually exclusive sets of alternatives. Typical examples include: aspect ratio; alternative subject choices based
5 on the minimal crop rectangle and maximal crop limits of various groupings of areas of interest; horizon placement alternatives (bottom third line or top third line); point of interest placement (at each of the four "one-third" intersection points or, for elongated items, along one of
10 the four "one-third" lines); and preferred edge placements for top, bottom, left and right edges.

For each edge there is a constraint set consisting of alternative ranges of distances that are acceptable on the
15 basis of an edge quality metric.

The examples given above are all "hard" constraints. That is, the condition must be met, and there is no gradual penalty involved in deviating from the condition. In many
20 cases it is desirable to implement constraints as "soft", that is, incurring an increasing penalty the further away the solution moves from the local optimum. An example is that positioning a horizon line exactly on the one-third line is better implemented in a way that allows placement
25 a little way off the precise one-third position, but penalises increasing distance from the desired one-third location.

The optimisation problem can easily be set to include
30 this. For example by changing the condition:

$$x = 1/3$$

to

$$x + e1 - e2 = 1/3$$

5

where $e1$ and $e2$ are positive penalty terms which contribute to an overall penalty function to be optimised, typically as a weighted sum of contributions such as...

10 $penalty = c1e1 + c2e2 + \dots$

The next step is to pick 94 the next combination of constraints, one alternative from each constraint set. Many combinations can be immediately excluded or
15 simplified as they are precluded by some of the other constraints. For example, choice of a particular cropping limits alternative will limit which points of interest can be considered as some may be outside those cropping limits.

20

The optimal crop candidate for the current constraints can then be determined 95. The constraints we have set up are combinations of simple linear conditions. These can be effectively solved by linear programming methods which
25 find the location for top, bottom, left and right boundaries of the crop which meet all the hard constraints and satisfy the soft constraints in such a way that the overall penalty is minimised.

30 Depending on the precise set of constraint combination

being solved, there may be a number of situations. Ideally, there is a single optimal solution.

However, there may be no solution. This would be the case
5 if some constraints were contradictory. For example, if
there are two points of interest A and B where A is to the
left of B, a combination of constraints that attempts to
place A near the right-hand one-third line and B near the
left-hand one-third line, then there is clearly no
10 solution. The method in step 94 of selecting sets of
constraints to solve should ideally be smart enough to
eliminate these situations.

There may be multiple solutions of equally low penalty
15 score. In this case we have a number of alternatives. One
is to pick a solution at random within the space of
multiple solutions. Another is to tighten the constraints,
for example by turning one or more soft constraints into
hard constraints. Optionally, in step 96 it is possible to
20 use a richer evaluation metric to generate a set of
alternatives within the space of equally acceptable
solutions and select these on the basis of the refined
evaluation metric. This optional step may, for example, be
a "generate and test" method. Many variants of this are
25 possible.

The linear solution is a practical method that will work
well for automated cropping because constraint sets can be
formed that represent alternative plausible choices.
30 Treating each combination independently and hence finding

different locally optimal solutions is a useful way of generating good alternatives for a user. Non-linear optimisation methods frequently suffer from problems with locally optimal solutions being confused for globally optimal solutions. An improved understanding of the search space allows this technique to circumvent such problems in a relatively intelligent manner.

Although linear programming is one method that may be used in step 95, it does impose limitations on the way the constraints are defined. Other optimisation techniques could be used within the same basic framework of local optimisation within a subspace defined by the choice of constraints from constraint sets.

If all possible constraints and evaluation criteria are encoded as hard or soft conditions which can be optimised in step 95, then step 96 may be bypassed. However, if step 95 is achieved by linear programming then some of the constraints may be poorly approximated or omitted. A more accurate evaluation of the solution generated by step 95 can be obtained afterwards in step 96. A more refined implementation might use the approximate solution from step 95 as the start point for a "generate and test" based local optimisation using the more detailed evaluation metrics.

An example of constraints that can be only approximated with a linear representation, is the edge quality metrics. The true edge quality can only really be assessed when the

limits of the edge are known. For example, the true relative quality of alternative left edge locations is dependent on the top and bottom limits. A narrower choice of top and bottom may exclude features in the image that would otherwise adversely affect the left edge quality. This type of interdependency cannot be modelled with a linear system. The best that can be done is that within step 93, having selected minimum and maximum cropping limits, the edge quality metrics are recalculated using, for example, the maximum cropping limits, to generate a specific set of alternative constraints for each edge constraint set. These specific constraints are used while the particular choice of cropping limits is being considered.

15

Another type of constraint that cannot be modelled linearly is one involving ratios of areas in the image, for example, the relative area of a boring region within the crop boundary. Clearly, this will be non-linear as the area is a multiplication of terms involving the horizontal and vertical crop locations.

20

In any event, once the crop candidate has been evaluated, this is recorded 97, along with its penalty score.

25

Then a test 98 is performed as to whether or not all constraint set combinations have been tested. If not 99, the flowchart loops back to step 94. If so, then there are many possibilities for deciding what to do with the results. Figure 7 shows just one example, in which the

30

crop candidates are sorted in order of penalty score 101, and then a number of these having the lowest penalty scores are presented 102 to a user of the system.

5 In addition to selecting from a set of alternatives, a user may wish to suggest improvements to the selected alternative. This could be achieved by simple commands such as "more" or "less". In such cases the system could define a new minimum crop rectangle and a new maximum crop
10 limits, based on the selected crop and the alternatives which the user rejected. For example, if the user requested "more" then the selected crop becomes the new minimum crop rectangle and the smallest non-selected crop which exceeds the size of the selected crop becomes the
15 new maximum crop limit. The system can then re-consider alternative edge placements within these limits to generate a new set of alternatives to present to the user. Repeated use of this form of interaction can allow the user to interactively explore the space of alternatives
20 which best meet the system's criteria for compositionally acceptable crops.

An example of another approach would be to ensure that one from each of the possible alternative crop limits was
25 used.

In some applications of the invention, a user may not be involved, and the crop may be fully automatic.

30 Figure 8 is a general example of a "generate and test"

method, in which certain constraints are defined as hard, equality constraints and are used to define a crop candidate without any attempt at optimisation. Each hard constraint is a local optimum location for a particular
5 criterion.

For ease of comparison, steps in Figure 8 that correspond with those in Figure 7 are given a reference numeral incremented by 100.

10

Once the relevant constraint sets for the features have been determined 192, these are enumerated 193A, and a group of constraint sets is selected 193B as "driving constraint sets". These constraint sets are such that when
15 groups of constraints are formed, one from each driving constraint set, a crop candidate is fully specified.

A simple example is for the group of driving constraints to consist of the constraint sets for top, bottom, left
20 and right locations, where each of these constraints is a candidate edge position determined from the edge quality metric.

So, for example, the left edge constraint set might be:

25

Left = 5

Left = 38

Left = 150

30 Analogous constraint sets would exist for possible Right,

Top and Bottom candidates.

In the example discussed above, the driving constraints are edge constraints such as these, combined with the
5 cropping limits for various combinations of areas of interest (i.e. minimum crop rectangles and cropping limits).

An alternative group of driving constraint sets might be
10 three enumerated edge location constraint sets and an aspect ratio constraint set.

The driving constraint sets determine which out of all possible crop rectangles are "generated". In the least
15 intelligent possible implementation, all possible left, right, top and bottom locations are generated. This, however, increases computational effort.

In step 194, the next combination of driving constraints
20 is selected, one from each driving constraint set.

The determination of the current crop candidate in step 195 is trivial, as there is only one possible solution to the driving constraints by definition.

25

In step 196, all soft constraints are evaluated. Here, an unspecified collection of soft evaluation criteria are combined.

30 Once the crop candidate has been evaluated, this is

recorded 197, along with its penalty score.

Then a test 198 is performed as to whether or not all driving constraint set combinations have been tested. If
5 not 199, the flowchart loops back to step 194. The flowchart only loops round the driving constraints as these fully specify the candidate crop rectangles to be considered.

10 If all driving constraint sets have been evaluated, then there are many possibilities for deciding what to do with the results. Figure 8 shows just one example, in which the crop candidates are sorted in order of penalty score 201, and then a number of these having the lowest penalty
15 scores are presented 202 to a user of the system.

The invention provides various advantages in automating or semi-automating the capture of cropped images with improved selection and composition of subject matter. The
20 invention also reduces the effort on the part of the photographer to take photographs with good and appropriate composition, thereby matching the potential of an electronic camera to capture a large number of images quickly.

25

Claims

1. An electronic image processing device for cropping an electronic image, comprising an image processing means,
5 the image processing means including an electronic processor and firmware and/or software for processing the electronic image, wherein the device is adapted to:
 - a) process the electronic image to identify one or more
10 features relevant to the composition of the electronic image, each such feature occupying a sub-region of the electronic image;
 - b) select one or more compositional rules from a set of
15 predetermined compositional rules, based on the relevance of the compositional rule(s) to one or more of the identified features; and
 - c) determine one or more suitable crop boundaries by
20 applying one or more of the selected compositional rules.
2. An electronic camera system for capturing an electronic image of a scene, comprising, a detector array for capturing the electronic image, an optical imaging
25 system arranged to image the scene onto the detector array, and an image processing device for cropping the captured image, wherein the image processing means is as claimed in Claim 1.
- 30 3. An electronic camera as claimed in Claim 2, in which

the camera includes means by which a user of the camera may tag one or more features relevant to the composition of the electronic image, said tagged feature(s) then being associated with a compositional rule that includes said
5 tagged feature(s) in the cropped image.

4. A method of using an electronic image processing device for cropping an electronic image, the image processing device comprising an image processing means,
10 the image processing means including an electronic processor and firmware and/or software for processing the electronic image, wherein the method comprises the steps of using the image processing means to:

- 15 i) process the electronic image to identify one or more features relevant to the composition of the electronic image each such feature occupying a sub-region of the electronic image;
- 20 ii) select one or more compositional rules from a set of predetermined compositional rules, based on the relevance of the compositional rule(s) to one or more of the identified features; and
- 25 iii) determine one or more suitable crop boundaries by applying one or more of the selected compositional rules.

5. A method as claimed in Claim 4, in which each identified feature has one or more compositionally
30 significant properties from amongst a plurality of

different predetermined compositional properties, and one or more compositional rules are selected from a set of predetermined compositional rules, based on the relevance of the compositional rule(s) to the compositionally
5 significant properties of one or more of the identified features.

6. A method as claimed in Claim 4 or Claim 5, in which step iii) comprises the steps of:

10

iv) generating a plurality of alternative candidate crop boundaries;

15 v) generating a measure of composition quality for each alternative candidate crop boundary by using a set of picture composition metrics; and

20 vi) selecting as output a reduced number of crop candidates having a relatively high measure of composition quality.

7. A method as in claimed in Claim 4 or Claim 5, in which step (iii) comprises the steps of:

25 vii) implementing the compositional rules as constraints that may be satisfied to a greater or lesser degree, each constraint having an associated cost function that increases the less well the constraint is satisfied.

30 viii) defining an overall cost metric as a function of

crop co-ordinates in the image, by forming a combination of the separate cost functions associated with each individual constraint;

5 ix) applying an optimisation method to find one or more best crop locations by finding minima in the overall cost metric; and

x) selecting as output a reduced number of crop candidates
10 for said best crop locations.

8. A method as claimed in Claim 6 or Claim 7, in which only one crop candidate is selected.

15 9. A method as claimed in any of Claims 4 to 8, in which the crop boundary of the cropped image has a plurality of crop edges, and the steps of the method are performed separately for each crop edge to generate the crop boundary.

20

10. A method as claimed in any of Claims 4 to 9, in which:

- step (i) includes identifying features which constitute
25 non-overlapping segmented regions of interest;

- step (iii) is achieved by first selecting alternative divisions of said identified regions of interest into those which will be included by the crop boundary and
30 those which will be excluded by the crop boundary;

- each such alternative division of regions of interest is used to determine an inner crop boundary limit and an outer crop boundary limit such that the inner crop boundary limit is the smallest boundary which
5 circumscribes the regions of interest to be included and the outer crop boundary limit is the largest boundary which excludes the regions of interest to be excluded; and

- said one or more suitable crop boundaries are determined
10 between the limits of the inner and outer crop boundary limits for each of said alternative divisions of regions of interest.

11. A method as in claimed in Claim 10 in which said
15 boundary has a plurality of edges and one or more suitable crop boundaries are determined by evaluating an edge quality metric for each possible edge placement between the inner and outer crop boundary limits, the selection of best crop boundary being at least in part dependent on the
20 separate edge quality metrics of each of the boundary's edges.

12. A method as claimed in any of Claims 4 to 11, in which one or more of said suitable crop boundaries are
25 presented to a user of the device for manual selection by the user for cropping of the electronic image.

13. A method as claimed any Claims 4 to 12, in which in
step (i) blank or uninteresting areas are detected as some
30 of the features relevant to the composition and in steps

(ii) and (iii) one or more of the compositional rules cause the image to be cropped according to the arrangement of said blank or uninteresting areas relative to other identified features relevant to the composition of the
5 electronic image.

14. A method as claimed in any of Claims 4 to 13, in which the image is a colour image and at least one feature is identified by segmenting the image into areas of
10 similar colour and/or texture.

15. A method as claimed in any of Claims 4 to 14, in which in step (i) features relevant to the composition of the electronic image are identified by a process of:

15

- segmenting the image into regions;

- denoting certain regions as more salient than others;

20 - grouping salient regions into larger regions separated by relatively non-salient regions;

- identifying said groups of salient regions as features with a region of interest property.

25

16. A method as claimed in Claim 15, in which the image is segmented into regions by segmenting the image into homogenous sub-areas, the measure of homogeneity being based on some combination of colour, intensity and
30 texture.

17. A method as claimed in Claim 15, in which certain regions are denoted as more salient than others by allocating high salience to regions on the basis of some combination of:

- relative unusualness of the colour, intensity or texture of the region to other adjacent regions; and/or

10 - relative unusualness of the colour, intensity or texture of the region relative to a substantial portion of the image.

18. A method as claimed in any of Claims 4 to 17, in which the compositional rule for features identified at the centre of the image is to include such features in the cropped image.

19. A method as claimed in any of claims 4 to 17, further comprising a step (0) prior to step (i) of determining, independently of capture of the electronic image, specific features or objects of interest, and wherein step (i) further comprises attempted identification of said specific features or objects of interest and designation of successfully identified specific features or objects of interest as features relevant to the composition of the electronic image.

20. A method as claimed in claim 19, wherein step (0) comprises identification of specific features or objects

prior to capture of the electronic image.

21. A method as claimed in claim 19, wherein step (0)
comprises provision of one or more identifier tags and of
5 rules for identification of identifier tags, and wherein
step (i) comprises use of said rules for identification of
said identifier tags.

22. A method as claimed in claim 21, wherein said
10 identifier tags are infra-red transmitters.

23. An electronic imaging processing device for cropping
an electronic image, substantially as herein described,
with reference to or as shown in the accompanying
15 drawings.

24. An electronic camera system for capturing an
electronic image of a scene, substantially as herein
described, with reference to or as shown in the
20 accompanying drawings.

25. A method of using an electronic image processing
device for cropping an electronic image, substantially as
herein described, with reference to or as shown in the
25 accompanying drawings.



Application No: GB 0031423.7
Claims searched: 1-25

Examiner: Matthew Males
Date of search: 21 June 2001

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.S): H4F (FFE FESK FESX FGXD)

Int Cl (Ed.7): H04N 1/387, 1/393, 5/262

Other: Online databases: WPI, EPODOC, JAPIO, INSPEC

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	EP 0824246 A2 XEROX - whole document but see abstract.	1, 4 at least
X	WO 99/09887 A1 QUALIA COMPUTING - whole document but see page 7, line 13-page 8, line 16.	1, 4 at least

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

THIS PAGE BLANK (USPTO)